EDGE DETECTION AND SHARPENING PROCESS FOR AN IMAGE

Background of the Invention

This invention relates to image enhancement during image processing. More particularly, the present invention relates to a method and an apparatus for sharpening edges of an image.

Monochromatic imaging systems seek to provide sharp edges and good background separation. When copying, printing or reproducing a document or image using a monochromatic imaging system, it is usually desirable to smooth photographic regions and to sharpen text regions. Text is sharpened by accentuating a contrast between a bordering dark region of text and an adjacent light background. Most imaging systems process images by first separating photographic regions from text regions, by means of a process called auto-separation, and then performing separate enhancement algorithms on the text regions and the photographic regions that have been identified by the auto-separation.

In processing black and white images (i.e. "monochromatic images"), systems treat image data as an array of discrete gray level pixels ("picture elements"). Each pixel is associated with a position in the image and an 8-bit digital value that encodes an intensity value in the range from 0 to 255. Each intensity value represents a gray level of the monochromatic image, ranging from absolute black (0) to absolute white (255). Alternately, the image data can be represented as a 12-bit digital value, comprising 4096 possible intensity values, ranging from 0 (black) to 4095 (white). In color images, the color data for the image may also be encoded in intensity values for the pixels. However with color images, each pixel is encoded by three intensity values, ranging from 0 to 255, or from 0 to 4095, which combine to define an output color for the pixel.

When a scanning device scans an image to capture image data, the scanning device corrupts the image and produces a blurring effect at edges within the image. An "edge" is an interface or boundary between two distinct regions in the image, such black text and a lighter background next to the text. This blurring effect, known as "scanner

fringe", occurs due to the scanner's inability to optically isolate one pixel in the image at a time. This effect is illustrated in Figure 1, which shows a typical scan measurement for a single pixel. A measurement in intensity of a high-intensity pixel 6 "ripples" to areas 7, 8 surrounding the pixel. When the scanner attempts to measure the intensity of a pixel in the image, neighboring pixels influence and affect the measurement. As a result, there is not a sharp transition between regions of low intensity pixels and regions of high intensity pixels. Rather, there is a fuzzy and blurred transition. For example, a low intensity pixel that borders a region of high intensity appears to the scanner to have a higher intensity, due to the influence of neighboring pixels in the high intensity region.

A second problem arises when the image is printed on a print medium with a half-tone printer. The half-tone printer operates in a binary mode, where the printer either prints a dot at a pixel position or leaves the pixel position blank. As such, the grey scale intensity values (which may range for instance from 0 to 255) must be mapped to one of two values. One of the values corresponds to a value for printing a dot and the other value is for not printing a dot. To effect this mapping, a print threshold is typically used. If the measured intensity level of a pixel is less than the print threshold, the output intensity value of the pixel is converted to 0 (absolute black) and the printer places a dot at the pixel position in the output image. If the measured intensity value of the pixel is greater than the print threshold, the output intensity value is converted to 255 (absolute white), and no dot is printed at the pixel location.

Error diffusion is a technique commonly used during this half-tone printing process for reducing error and accurately reproducing an average intensity level for the image. Error diffusion examines each pixel, determines an error value for the pixel and forwards the error to a selected collection of neighboring pixels in accordance with a weighting scheme. The error is the difference between the gray level (0-255) pixel value of the input digital image and the output intensity value (either 255 or 0) of what is printed. In this manner, the error resulting from the conversion is diffused to surrounding pixels, and the average intensity (the gray level) of the original image is maintained in the converted image.

Although error diffusion may help to decrease image corruption during half-tone printing, the printer still corrupts the image at the edges of the image. Edges become blurred as the printer attempts to match the gray level at a transition between a region of low intensity pixels and a region of high intensity pixels, such as the border between text and a lighter background, as commonly found at edges of an image. The error diffuser may diffuse error to a light side of an edge region, and place ink where no dot should be placed.

Summary of the Invention

The present invention addresses the above-described problems through an improved method and device for sharpening edges in an image. The invention corrects for a distortion of the image signal at an edge that occurs during scanning and printing of the image. This is accomplished by applying a sharpening filter to detected edges of an image only. The filter detects the edges of an image and sharpens the edges according to the amount of distortion in the image signal. Unlike the classical approach, this filter operates on the fly, without an auto separation process.

Rather than classifying regions of an image as text or photo, the invention scans across the whole image on a pixel by pixel basis, and finds areas of high contrast that define an edge region. These regions include text printed on light backgrounds and photo details that are blurred by the scanner. The filter changes the gray response levels of the image signal in a way tuned to the print engine so that details and borders are sharper. Edges are enhanced by altering the manner in which a halftone print engine places ink at an edge of a printed image. Sharpening accentuates the difference between the two sides of an edge region. This invention improves both photo and text quality by only adjusting a small percentage of the image data. Smooth transition areas in photos are unaffected, but edges of objects within a photo are enhanced.

According to one embodiment, the invention manipulates the image signal in an edge region to increase the intensity level of output pixels on the light side of the transition and to decrease the intensity level of output pixels on the dark side of the

transition. This effect increases contrast in an edge region and prevent the printer from placing ink on the light side of the transition.

According to another embodiment, the invention manipulates the image signal in an edge region to increase contrast and provide detail to the edge. This is accomplished by applying a both a positive and a negative gain to the image signal in both regions that border an edge. While the portion of dark pixels that directly border the edge are darkened, a portion of the pixels in the dark region that are farther from the edge are lightened. In addition, the portion of light pixels that directly border the edge are lightened, and the portion of the pixels in the dark region that are farther from the edge are darkened. In this embodiment, the edge is further emphasized by providing greater contrast and detail to the edge region.

According to another embodiment, the invention provides a method of enhancing the edges of an image in order to compensate for distortion that occurs during scanning of the image.

According to yet another embodiment, the invention provides a method of enhancing the edges of an image in order to compensate for distortion that occurs during printing of the image.

Brief Description of the Drawings

An illustrative embodiment of the present invention will be described below relative to the following drawings.

Figure 1 is a graph of pixel intensity versus position for an isolated pixel illustrating the effect of "scanner fringe" in a conventional system.

Figure 2 is a diagram of a scan-to-print system suitable for practicing the illustrative embodiment of the present invention.

Figure 3A is a graph of pixel intensity versus position for input image data at an edge region.

Figure 3B is a graph of pixel intensity versus position for image data of Figure 3A that is corrupted by a scanner at an edge region.

Figure 3C is a graph of pixel intensity versus position for image data of Figure 3A that is corrupted by the printer at an edge region.

Figure 4A is one embodiment of a filter pixel window for a horizontal implementation.

Figure 4B is a second embodiment of the filter pixel window for a vertical implementation.

Figure 5A illustrates one arrangement for pixel window contexts.

Figure 5B illustrates a second arrangement for pixel window contexts.

Figure 6 is a block diagram illustrating a filtering process for a horizontal application of the filter.

Figure 7A is a graph of the image signal in an edge region after application of the filter in one embodiment of the present invention.

Figure 7B shows a detail of an edge to which the filter of the embodiment depicted in Figure 7A has been applied.

Figure 8A is a graph of the image signal in an edge region after application of the filter in an alternate embodiment of the invention.

Figure 8B is a block diagram detailing the step of assigning a gain value for the embodiment of Figure 8A.

Figure 8C shows detail of an edge to which the filter illustrated in Figure 8A has been applied.

Detailed Description of the Invention

The illustrative embodiment of the present invention provides a method and device for sharpening edges of an image during image processing. The illustrative embodiment will be described below relative to an implementation in a scan-to-print system for generating a copy of an original image. Nevertheless, those skilled in the art will appreciate that the present invention may also be implemented on other types of image forming systems, such as a fax machine, a personal computer or a photocopier.

The illustrative embodiment enhances the quality of a printed image by sharpening the edges of the image. The present invention allows a user to adjust the

degree of sharpening according to the amount of distortion in the image. The present invention operates in real time, to improve image quality as image data is processed by an imaging system.

Figure 2 illustrates a scan-to-print image forming system 9 suitable for practicing the illustrative embodiment of the present invention. An image-capturing device, such as a scanner 1, scans an original image 2 and converts image data to an electronic signal. The signal is transmitted to the signal processor 3, which processes the image signal. The image signal is then sent to the printer 4, which produces a copy 5 of the original image on a print medium. An image forming system includes image development and transfer systems comprising an assemblage of operatively associated image forming elements for capturing image data related to an input image, generating a latent image corresponding to the input image, depositing a latent image onto a receiving member, developing the image, and then transferring the developed image onto a print medium.

Figures 3A, 3B, and 3C illustrate the two distortions of an image signal. Figure 3A is a graph of a typical input data signal 10 at an edge 11. A sharp transition exists between a region of high intensity pixels 12 and a region of low intensity pixels 13 in at an edge in an image. Figure 3B shows the corruption of the edge that occurs in the image signal 10 during the scan process. Rather than a sharp transition at the edge, as shown in Figure 3A, the data signal 10 from the scanner undergoes a more gradual transition at an edge 11. The scanner 1 tends to merge the intensity levels of the two regions, rather than distinguish two discrete intensity levels. The width of the distortion W_1 is two pixels wide, such that the distortion extends one pixel on either side of the edge 11. Figure 3C shows a second similar distortion that occurs during the printing process. The printed output from the printer 4 also merges the two regions on either side of the edge 11 as the error from the half-tone printer 4 is carried forward to an edge region by an error diffuser in the printer 4. The width of the printer distortion W_2 is one pixel wide. The combined effect of the two distortions on the image signal is a fuzzy edge and an indistinct transition between the two intensity regions of the image.

In order to correct the distortion that is illustrated in figures 3B and 3C, a filter, contained within the signal processor 3 is applied to the image signal during image processing. By manipulating the signal before the error diffusion process, the filter essentially performs the inverse of the corruption on the image signal.

The filter may be applied either one dimensionally or two dimensionally, as necessary. However, for most images and character sets, the filter will only necessitate a horizontal application in order to provide acceptable results. A horizontal application of the filter sharpens non-horizontal (i.e. vertical) edges of the image. A vertical application of the filter, from top to bottom, is also possible and sharpens non-vertical (i.e. horizontal) edges of the image.

The filter detects edges by analyzing each pixel in the image in the context of the surrounding image area. As illustrated in figure 4A, the filter first creates a pixel window for every pair of adjacent pixels in the image in order to detect edges. The pixel window 20 for a pair of pixels 21, 22 being processed includes a left pixel 21, a right pixel 22, a left context 23, and a right context 24. This configuration is applied in horizontal scanning and filtering.

Figure 4B shows an alternate pixel window configuration that is used in vertical scanning to sharpen horizontal edges. In this embodiment, the pixels being processed are one on top of the other. This pixel window 20' includes the top pixel 21', and the bottom pixel 22', a top context 23' and a bottom context 24'.

In one embodiment, shown in figure 5A, the left context 23 and the right context 24 are four pixels wide, and are contiguous with the left pixel 21 and the right pixel 22, respectively. In an alternate embodiment, shown in figure 5B, the left context 23 overlaps and includes the left pixel 21, and the right context 24 overlaps and includes the right pixel 22. The size of the surrounding contexts, as well as the distance in pixels between the contexts and the pixels is dependent on the individual scanner or printer that is in use. However, in the preferred embodiment, the left, right, top and

bottom contexts are four pixels wide and contiguous with the processed pixel pair in order to provide accurate compensation for the edge corruption.

Figure 6 demonstrates the filtering process for a horizontal implementation of the filter. After defining the pixel window for the processed pixel pair (step 100), the filter detects whether the pixel pair form part of an edge in the image by calculating the average intensity value of the left context (step 101) and the average intensity value of the right context (step 102). The average intensity values are computed by measuring the intensity of each pixel within the context, finding the sum of these intensities, and dividing by the number of pixels in the context. The filter then computes the difference between the average intensity value of the left context and the average intensity value of the right context (step 103). The filter squares the difference value to create a positive magnitude value (step 104). When an edge exists in the image, there is a considerable disparity between the average intensity of the left context and the average intensity of the right context. The magnitude value is compared to a threshold value to determine whether an edge exists or not (step 105). If the difference between the left context and the right context is too low to be an edge region, a gain value is set to zero (step 106), and no adjustment of the pixel values is made. If, however, an edge exists at that location in the image, an appropriate gain value is assigned, depending on the degree of sharpening required (step 107). The step of assigning a gain value will be explained in further detail below. The filter stores the gain value (step 109) and the difference value (110) for use in subsequent calculations.

In a typical sharpening process, a difference between the two contexts that spans about 60 percent of the range of pixel values (0-255) indicates that an edge is present. For example, if a left and right pixel straddle a transition, and the left pixel is part of a dark region and the right pixel is part of a light region, the difference between the two intensity values of the left and right context is a relatively high number. The average intensity value of pixels in a dark region is usually 50 or less, and the average intensity value of pixels in a light region is typically 200 or more. Thus, a difference between the average intensity values in a pixel window of about 150 or more indicates the presence of an edge in the image.

On the other hand, a smooth transition in intensity over a region has a smaller difference value. The filter recognizes that no edge is present and, therefore, does not manipulate the output signal.

When the filter detects an edge, it adjusts the intensity of the processed pixels to compensate for the expected distortion of the image at that edge (step 108). In the illustrative embodiment, the intensity of each pixel is increased or decreased, according to the following formulas:

i) $X_{left} = X_{left} + (gain*difference) - (last\_computed\_gain* last\_computed\_difference)$

and

ii) $X_{right} = X_{right} - (gain*difference) + (last\_computed\_gain*last\_computed\_difference)$

The filter adjusts the intensity value of each pixel of the pixel pair (X_{left} and X_{right}) so that pixels in the light region of the edge are further lightened and pixels in the dark region of the edge are further darkened. For the left pixel ("X_{left}"), the intensity value is adjusted by adding the product of the gain value ("gain") and the difference value ("difference") corresponding to the currently processed pixel window to the measured intensity value of the pixel. The filter then subtracts the product of the stored gain value ("last_computed_gain") and the stored difference value ("last_computed_difference") for the pixel pair that was processed directly before the current pixel pair. These values are initially set as zero. For the right pixel, the intensity value is adjusted in a reciprocal fashion. The filter modifies the right pixel value (X_{right}") by subtracting the product of the gain value ("gain") and the difference value ("difference") and then adding the product of the stored difference value ("last_computed_difference") and the stored gain value ("last_computed_gain") from the previously processed pixel pair. In this manner, the average pixel intensity value at

an edge is maintained, while the difference between the light and dark pixels is intensified.

The filter then shifts the pixel window (step 111) by one pixel in order to process the next pixel pair, and adjust the output intensities, if necessary. In the case of horizontal filtering, the pixel window may be shifted from left to right. The right pixel of the pixel window becomes the new left pixel, the left pixel becomes a part of the left context, and the left and right contexts are adjusted accordingly. The filter continues to perform the same filtering process illustrated in Figure 6 for the new pixel pair, using the stored gain value and the stored difference value from the previous pixel pair. In the case of vertical filtering, the pixel window is shifted from the top of the page to the bottom of the page. In this embodiment, the bottom pixel becomes the top pixel, the top pixel becomes part of the top context and the top and bottom contexts are shifted accordingly. In many instances, each pixel that is part of an edge region will be adjusted twice: the first time as the right pixel, and the second time as the left pixel. When shifting, the filter will read only the newest pixel that is added to the pixel pair, which is the rightmost pixel of the right context in the case of horizontal filtering, and use the stored intensity values of the rest of the pixel window. Therefore, unmodified values are always used when computing the average intensity values.

The threshold value and the gain values are adjustable by the user, according to the desired degree of filtering. A first filter response, illustrated in Figure 7A, utilizes one threshold value in determining the gain value. When the filter detects an edge 211 (the magnitude value is greater than the threshold value), a negative gain value is assigned in the formula to adjust the pixel outputs (steps 107 and 108). In this first filter response, the filter applies a negative gain 214 to the image signal 210 in the low intensity (dark) side 213 of the edge 211 , and a positive gain 215 to the image signal 210 in the high intensity (light) side 212 of the edge 211 . The negative gain 214 decreases the intensity of the pixels in the low intensity region of the edge, while the positive gain 215 increases the intensity of the pixels in the high intensity region of the edge. In a preferred embodiment, the threshold for assigning a gain value is when the difference is greater than 140 (the magnitude is greater than 19600). The gain value is

set in the range of 0 to –1. In a preferred embodiment, the gain value assigned in Equations i) and ii) when an edge is detected is –0.75 for this first filter response.

Figure 7B illustrates a detail of an edge 211 to which this filter response has been applied. The application of this filter response darkens one portion 214 of the edge in the region of low intensity pixels 213 while simultaneously lightening 215 an adjacent portion of the edge that is in the region of high intensity pixels 212 . This increases the contrast at the edge between the two distinct regions, while maintaining a substantially constant average intensity level for the image. In this manner, the filter clarifies the fuzzy transition in the signal at this point. When the signal is error diffused, the half-tone printer will tend to place dots on one side of the transition, and prevents ink from being placed in the light region of the edge.

Figure 8A shows an alternate filter response. This second response provides added detail and further contrast to the edge region by using two different threshold values when assigning a gain value in steps 105, 106 and 107 of Figure 6. Figure 8B is a flow chart illustrating the step of assigning a gain value in this response. In this response, a positive gain value, between 0.25 and 0.75 is assigned if the magnitude is greater than a first threshold value but less than a second threshold value. In this case, the first threshold is usually 40-50 percent of the range of the pixel values, and the second threshold is usually 60 percent of the range of the pixel values. A suitable first threshold value is about 10,500 ($[256*40\%]^2$), and a suitable second threshold value is about 23,600 ($[256*60\%]^2$). If the magnitude is greater than or equal to the second threshold value, the filter assigns a negative gain value between 0 and –1. When applying the formula in step 108, a first negative gain 314 similar to the negative gain 214 of figures 7B and 7B is applied to the region of low intensity that immediately borders the edge 311. A first positive gain 316 is also applied to slightly increase the intensity of the low-intensity pixels that are in an intermediate zone between the edge and the principal low intensity region of the image (i.e. the low intensity region that does not directly border the edge). On the high intensity side of the edge, a second positive gain 315 , similar to the positive gain 215 of figures 7A and 7B, serves to increase the intensity of pixels in the high intensity region that directly border the edge.

A second negative gain 317 serves to slightly decrease the intensity of pixels in the high intensity region that are in an intermediate zone between the edge and the principal high intensity region of the image.

The effect of the second filter response at an edge is shown in Figure 8C. This enhancement to the edge provides greater detail and heightened contrast to an edge region. The added detail emphasizes the edge, and further highlights the contrast between the two regions in the image while maintaining an average intensity level on both sides of the edge. In addition to clarifying a fuzzy transition, the filter accentuates an edge portion, thereby enhancing overall image quality. The second filter response regulates the signal, so that the half-tone print engine will cluster dots on the appropriate darker side of the transition, and leave the immediately bordering lighter side of the transition white.

The filter is flexible according to the desired amount of filtering, and may be tuned to a particular printer and scanner combination. In addition, the filter may be modified so as to compensate for distortion from the scanner only, or to compensate for distortion from the printer only. The assigned gain values, and the threshold values for assigning the gain values are adjustable so as to provide a specific response. In addition, the contexts in the pixel window may be modified to different sizes and shapes. For example, a printer that is unable to handle a large number of overly sharpened transitions quickly may require a larger context size. This prevents the edge filter from being activated to adjust the pixel outputs too frequently. The filter may be tuned through a variety of different settings, and an appropriate combination of settings may be realized by one of ordinary skill in the art through routine experimentation.

While the invention has been described for processing monochrome images, it is recognized that the filter is also capable of sharpening color images in an analogous manner. By running the filter on the image's luminance channel, edges of color images are sharpened as well. Alternatively, the filter runs on the individual output color planes and sharpens accordingly.

While the invention has been described according to particular embodiment, it will be recognized by one of ordinary skill in the art that modifications and alterations may be made without departing from the spirit and scope of the invention. The specification and drawing are intended to be exemplary, rather than restrictive. Therefore, the scope of the invention is to be determined by reference to the appended claims.